arXiv:1606.05988v1 [stat.ME] 20 Jun 2016

Continuum directions for supervised dimension reduction

Sungkyu Jung

Department of Statistics, University of Pittsburgh
Pittsburgh, PA 15222, U.S.A.
e-mail: sungkyu@pitt.edu

Abstract: We consider dimension reduction of multivariate data under the existence of various types of auxiliary information. We propose a criterion that provides a series of orthogonal directional vectors, that form a basis for dimension reduction. The proposed method can be thought of as an extension from the continuum regression, and the resulting basis is called continuum directions. We show that these directions continuously bridge the principal component, mean difference and linear discriminant directions, thus ranging from unsupervised to fully supervised dimension reduction. With a presence of binary supervision data, the proposed directions can be directly used for a two-group classification. Numerical studies show that the proposed method works well in high-dimensional settings where the variance of the first principal component is much larger than the rest.

Primary 60K35, 60K35; secondary 60K35.
Keywords and phrases: continuum regression, dimension reduction, Fisher discriminant analysis, high dimension low sample size, maximum data piling, principal component analysis.

1. Introduction

In modern complex data, it becomes increasingly common that multiple related data sets are available. We consider the data situation where a supervised dimension reduction is naturally considered. Two types of data are collected on a same set of subjects: a data set of primary interest $\mathcal{X}$ and an auxiliary data set $\mathcal{Y}$. The goal of supervised dimension reduction is to delineate major signals in $\mathcal{X}$, dependent to $\mathcal{Y}$. Relevant application areas include genomics (genetic studies collect both gene expression and SNP data (Li et al., 2015a)) and finance data (stocks as $\mathcal{X}$ in relation to characteristics $\mathcal{Y}$ of each stock: size, value, momentum and volatility (Connor, Hagmann and Linton, 2012)), and batch effect adjustments (Lee, Dobbin and Ahn, 2014).

There has been a number of work in dealing with the multi-source data situation. Lock et al. (2013) developed JIVE to separate joint variation from individual variations. Large-scale correlation studies can identify millions of pairwise variable associations, either between genomic data types (Gilad, Rifkin and Pritchard, 2008), or for any data types through Canonical Correlation Analysis (CCA, Hotelling, 1936) and multiple CCA (Witten and Tibshirani, 2009). Graphical models (Yang et al., 2014) can be efficient in summarizing statistical

dependence. Most of these methods, however, do not provide supervised dimension reduction of a particular data set $\mathcal{X}$, since auxiliary data sets assume an equal role.

In contrast, reduced-rank regression (RRR, Izenman, 1975; Tso, 1981) and envelop models (Cook, Li and Chiaromonte, 2010) provide sufficient dimension reduction (Cook and Ni, 2005) for regression problems. Variants of Principal Component Analysis (PCA, Wold, Kettaneh and Tjessem, 1996; Fan, Liao and Wang, 2014; Di et al., 2009) has been proposed to incorporate auxillary information. Recently, Li et al. (2015b) proposed SupSVD, a supervised PCA that encompasses regular PCA to RRR. Our goal is similar to that of SupSVD, which extends RRR and envelop models, in that main and auxiliary data sets play different roles. We consider a basis (or subspace) recovery to extract the part of main data set which is relevant to the auxiliary data set. Unlike SupSVD, which provides a fully supervised dimension reduction, we seek a unified framework that covers a wide spectrum from fully-supervised to unsupervised dimension reduction.

A potential drawback of fully supervised dimension reduction as a preprocessing for further application of predictive modeling is a *double-dipping* problem: the same signal is considered both at dimension reduction and at classifiers. In high dimensional data situations, small signals can sway the whole analysis, often leading to a spurious finding that can not be replicated in subsequent studies. A regularized semi-supervised dimension reduction has a potential to mitigate the double-dipping problem.

In this work, we propose a semi-supervised basis learning for the primary data that covers a wide range of spectrum from supervised to unsupervised dimension reduction. A meta-parameter $\gamma \in [0, \infty)$ is introduced to control the degrees of supervision. The spectrum of dimension reduction given by different γ is better understood when there exists a single binary supervision. In such a special case, the directional vectors of the basis continuously bridge the principal component direction, mean difference and Fisher's linear discriminant directions.

The proposed method was motivated by the continuum regression (Stone and Brooks, 1990), regressors ranging from the ordinary least square to the principal component regression. In the context of regression, our primary data set is predictors while the auxiliary data are the response. The new basis proposed in this work, called *continuum directions*, can be used with multiple supervision data, consisting of either categorical or continuous variables.

We also pay a close attention to the high-dimension, low-sample-size situations (or the $p \gg n$ case), and give a new insight on the maximum data piling (MDP) direction w_{MDP}, proposed as a discriminant direction by Ahn and Marron (2010). In particular, we show that w_{MDP} is a special case of the proposed continuum direction, and if $p \gg n$, MDP is preferable to linear discriminant directions in terms of Fisher's original criterion for Fisher's linear discriminant analysis (LDA, Fisher, 1936).

As an application of the continuum directions, we endeavor to use the continuum directions in the binary classification problems. Recently, numerous efforts to improve classifications for high dimension, low sample size contexts, or

$p \gg n$, have been made. Linear classifiers such as LDA, the support vector machine (Vapnik, 1995; Cristianini and Shawe-Taylor, 2000) or distance weighted discrimination (Marron, Todd and Ahn, 2007; Qiao et al., 2009) often yield better classification than nonlinear methods, in high dimensional data analysis. A recent trend is sparse estimations. Bickel and Levina (2004) studied the independence rule, ignoring off-diagonal entries of S_W. Additionally, assuming sparsity of the population mean difference Fan and Fan (2008) proposed the features annealed independence rule (FAIR). Shao et al. (2011) considered a sparse LDA based on thresholding, while Wu et al. (2009) proposed another sparse LDA with an ℓ_1 constraint. Cai and Liu (2011) proposed the linear programming discriminant rule (LPD) for sparse estimation of the discriminant direction vector. The sparse LDA and LPD are designed to work well if their sparsity assumptions are satisfied. Sophisticated methods such as those of Wu et al. (2009) and Cai and Liu (2011) usually suffer from heavy computational cost. On the other hand, simpler methods such as FAIR and MDP are much faster. Our method, when applied to the binary classification problem, leads to analytic solutions, and the computation times are scalable.

2. Continuum directions

2.1. Motivation

To motivate the proposed directions for dimension reduction, we first analyze a special case where the supervision data consist of a binary variable. We discuss a few meaningful directions for such situations, viewed in terms of a two-group classification problem. These directions are special cases of the continuum directions, defined later in (6)

Let n_1 and n_2 be the numbers of observations in each group and $n = n_1 + n_2$. Denote $\{x_1, \ldots, x_{n_1}\}$ and $\{y_1, \ldots, y_{n_2}\}$ for the p-dimensional observations of the first and second group, respectively. In our study it is sufficient to keep the sample variance-covariances. Denote $S_W = \frac{1}{n}(\sum_{i=1}^{n_1}(x_i-\bar{x})(x_i-\bar{x})^T + \sum_{i=1}^{n_2}(y_i-\bar{y})(y_i-\bar{y})^T)$ for the within-group variance matrix, i.e. the estimated (pooled) common covariance, and $S_B = \frac{n_1 n_2}{(n_1+n_2)^2}(\bar{x}-\bar{y})(\bar{x}-\bar{y})^T$ for the between-group variance matrix. The total variance matrix is $S_T = \frac{1}{n}(\sum_{i=1}^{n_1}(x_i-\hat{\mu})(x_i-\hat{\mu})^T + \sum_{i=1}^{n_2}(y_i-\hat{\mu})(y_i-\hat{\mu})^T)$ with the common mean $\hat{\mu} = \frac{n_1\bar{x}+n_2\bar{y}}{n}$, and $S_T = S_W + S_B$.

Fisher's criterion for discriminant directions is to find a direction vector w such that, when data are projected onto w, the between-variance $w^T S_B w$ is maximized while the within-variance $w^T S_W w$ is minimized. That is, one wishes to find a maximum of

$$T(w) = \frac{w^T S_B w}{w^T S_W w}. \tag{1}$$

If S_W is non-singular, i.e. the data are not collinear and $p \leq n-2$, the solution is given by $w_{LDA} \propto S_W^{-1} d$, where $d = \bar{x} - \bar{y}$. It has been a common practice to extend the solution to the case $p > n-2$ using a generalized inverse, i.e.,

$$w_{LDA} \propto S_W^{-} d,$$

where A^- here stands for the Moore–Penrose pseudoinverse of square matrix A.

In retrospect, when $\text{rank}(S_W) < p$, Fisher's criterion is ill-posed since there are infinitely many w's satisfying $w^T S_W w = 0$. Any such w, which also satisfies $w^T S_B w > 0$, leads to $T(w) = \infty$. In fact, in such a situation, w_{LDA} is not a maximizer of T but merely a critical point of T. Ahn and Marron (2010) proposed a maximal data piling (MDP) direction w_{MDP} which maximizes the between-group variance $w^T S_B w$ subject to $w^T S_W w = 0$, and is

$$w_{MDP} \propto S_T^- d.$$

Note that w_{MDP} also maximizes a criterion

$$T_{MDP}(w) = \frac{w^T S_B w}{w^T S_T w}. \tag{2}$$

In the conventional case where $n \geq p$, the criteria (1) and (2) are equivalent up to a constant, and $w_{MDP} = w_{LDA}$. We discuss further in Section 2.4 that MDP is more preferable than LDA in the high-dimensional situations.

A widely used modification to Fisher's criterion is to shrink S_W towards an identity matrix, leading to

$$T_\alpha(w) = \frac{w^T S_B w}{w^T (S_W + \alpha I) w}, \quad \text{for some } \alpha \geq 0. \tag{3}$$

This approach has been understood as similar in flavor to ridge regression (Hastie, Tibshirani and Friedman (2009); Wu et al. (2009)). The solution of the above criterion is simply given by $w_\alpha^R \propto (S_W + \alpha I)^{-1} d$. A special case is in the limit $\alpha \to \infty$, where the solution w_∞^R becomes the direction of mean difference (MD) $w_{MD} \propto d$, which maximizes

$$T_{MD}(w) = w^T S_B w, \tag{4}$$

with a conventional constraint $w^T w = 1$.

In high dimensional data situations, utilizing the principal components is a natural and nonparametric way to filter out any redundant noise. Principal component analysis (PCA) reduces the dimension p to some low number p_0 so that the subspace formed by the first p_0 principal component directions contains maximal variation of the data among all other p_0 dimensional subspaces. In particular, the first principal component direction w_{PC1} maximizes the criterion for the first principal component direction,

$$T_{PCA}(w) = \frac{w^T S_T w}{w^T w}. \tag{5}$$

The important three directions of MDP, MD and PCA differ only in criteria maximized. With the constraint $w^T w = 1$, the criteria (2)–(5) are functions of total-variance $w^T S_T w$ and between-variance $w^T S_B w$. For the special binary supervision case, a generalized criterion that embraces all three methods is

$$T_\gamma(w) = (w^T S_B w)(w^T S_T w)^{\gamma - 1} \quad \text{subject to} \quad w^T w = 1, \tag{6}$$

where γ takes some value in $[0, \infty)$. The special cases are MDP as $\gamma \to 0$, MD at $\gamma = 1$, and PCA when $\gamma \to \infty$. The direction vector w_γ that maximizes T_γ is called the *continuum direction* for γ.

2.2. General Continuum directions

The continuum direction (6) defined for the special binary supervision is now generalized to incorporate any form of supervision.

Denote $X = [x_1, \ldots, x_n]$ for the $p \times n$ primary data matrix and Y for the $r \times n$ matrix with secondary information. The matrix Y contains the supervision information that can be binary, categorical, and continuous. Assuming for simplicity that X is centered, we write the total variance-covariance matrix of X by $S_T = n^{-1}XX^T$, and the Y-relevant variance-covariance matrix of X by $S_B = n^{-1}(XY^T)(XY^T)^T$. A completely *unsupervised* dimension reduction can be obtained by eigenvalue decomposition of S_T. On the contrary, a fully-supervised approach is to focus on the column space of S_B, corresponding to the mean difference direction when Y is binary. An extreme approach that nullifies the variation in X to maximize the signals in Y can be obtained by eigenvalue decomposition of $S_T^- S_B$. When Y is categorical, this reduces to the linear discriminant analysis.

Generalizing (6), the following approach encompasses the whole spectrum from the supervised to unsupervised dimension reduction. A meta-parameter $\gamma \in [0, \infty)$ controls the degree of supervision. For each γ, we obtain a basis $\{w_{(1)}, \ldots, w_{(\kappa)}\}$ for dimension reduction of X in a sequential fashion. In particular, given $w_{(1)}, \ldots, w_{(k)}$, the $(k+1)$th direction is defined by w maximizing

$$\begin{aligned} T_\gamma(w) &= (w^T S_B w)(w^T S_T w)^{\gamma - 1} \\ \text{subject to } \; & w^T w = 1 \;\text{ and } w^T S_T w_{(\ell)} = 0, \; \ell = 1, \ldots, k. \end{aligned} \tag{7}$$

The sequence of directions $\{w_{(\ell)} : \ell = 1, \ldots, \kappa\}$ for a given value of γ is then S_T-orthonormal to each other ($w_{(\ell)}^T w_{(\ell)} = 1, w_{(\ell)}^T S_T w_{(l)} = 0$ for $\ell \neq l$). An advantage of requiring S_T-orthogonality is that the resulting scores $z_{\ell,i} = x_i^T w_{(\ell)}$ are uncorrelated with $z_{l,i}$ for $\ell \neq l$. These are desirable if these scores are used for further analysis, such as a classification based on these scores.

In sequentially solving (7), choosing large γ provides nearly unsupervised solutions while $\gamma \approx 0$ yields an extremely supervised dimension reduction. The spectrum from unsupervised to supervised dimension reduction is illustrated in a real data example shown in Example 1. The spectrum can also be understood in an analogy to the special binary case (6), the discussion of which is deferred to the next subsection.

Example 1. We demonstrate the proposed method of dimension reduction for a real data set from a microarray study. This data set, described in detail in Bhattacharjee et al. (2001), contains $p = 2530$ genes (primary data) from $n = 56$

patients while the patients are labeled by four different lung cancer subtypes (supervision data).

The continuum directions can provide basis of dimension reduction, ranging from unsupervised ($\gamma \approx \infty$) to fully supervised ($\gamma \approx 0$) dimension reduction. In Fig. 1, the projected scores of the original data are plotted for four choices of γ.

A dimension reduction by PCA has been useful for this data set, since the four subtypes are visually separated by using the first few sample principal components. The principal component scores are similar to those plotted in the first panel of Fig. 1 when γ is large enough. On the other hand, a fully supervised dimension reduction given by the MDP directions, plotted in the bottom right panel, nullifies any variation in the primary data set. Specifically, all observations corresponding to the same subtype project to a single point, a feature due to the high dimensionality. Thus the projected score for $\gamma = 0$ contain only the information relevant to the supervision.

The continuum directions as functions of γ are generically continuous (shown later in Proposition 2), thus the projected scores are also continuous with respect to γ. The continuous transition of the scores from large γ to small γ in Fig. 1 is thus expected. The question of which value of γ to use in final dimension reduction depends on the purpose of analysis. In exploratory statistics of the data, several values of γ may be used to give a full exploration of the data set. If the dimension reduction is performed for regression or classification, a cross-validation can be used. We further investigate a cross-validatory choice of γ for a special two-group classification study in Section 3.

2.3. Relation to continuum regression

A special case of the proposed method, specifically (6) for the binary supervision, can be viewed as a special case of continuum regression (Stone and Brooks, 1990). The continuum regression leads to a series of regressors that bridges ordinary least squares, partial least squares and principal component regressions. In connection with the continuum directions for binary classification, ordinary least squares regression corresponds to LDA (or MDP in (2)), and partial least squares corresponds to mean difference. In particular, in the traditional case where $n > p$, it is widely known that w_{LDA} is identical to the coefficients of least squares regression, up to some constant (Hastie, Tibshirani and Friedman, 2009, p.135).

Some related work has shed light on the relationship between continuum regression and ridge regression (Sundberg, 1993; de Jong and Farebrother, 1994; Bjorkstrom and Sundberg, 1999). A similar relationship can be established for our case when S_B is of rank 1. For simplicity, we assume that the column space of S_B is spanned by the vector d. To find the continuum direction w_γ that maximizes $T_\gamma(w)$ in (7), differentiating the Lagrangian function $\log T_\gamma(w) - \lambda(w^T w - 1)$ with respect to w leads to the equation

$$\frac{S_B w}{w^T S_B w} + (\gamma - 1)\frac{S_T w}{w^T S_T w} - \lambda w = 0. \tag{8}$$

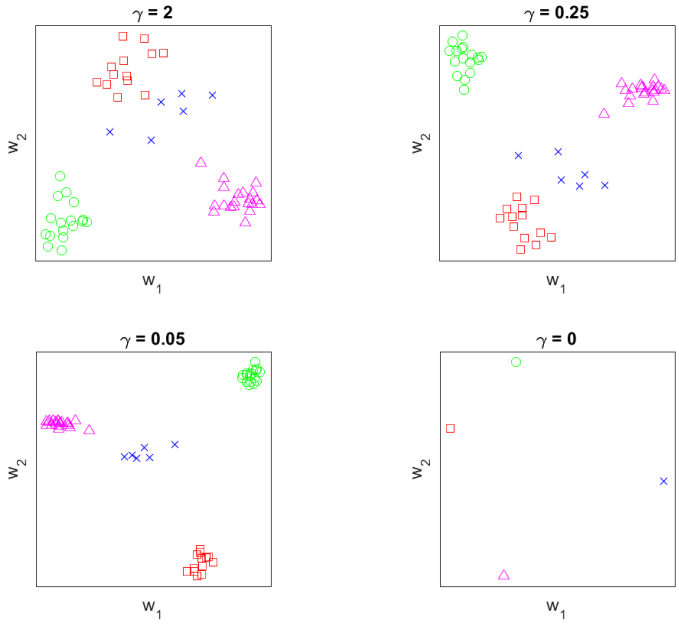

FIG 1. *Spectrums of supervised dimension reduction for the data set of Bhattacharjee et al. (2001). Shown are the projection scores to the first two continuum directions, for various values of* γ.

Left multiplication of w^T leads to $\lambda = \gamma$. A critical point of the preceding equation system gives the maximum of T_γ. Since $\frac{S_B w}{w^T S_B w} = \frac{dd^T w}{w^T dd^T w} = \frac{1}{d^T w} d$, one can further simplify the equation for a critical point

$$w \propto (S_T + \frac{\gamma}{1-\gamma} \frac{w^T S_T w}{w^T w} I_p)^- d = (S_T + \alpha I_p)^- d := w_\alpha^R.$$

For each $\gamma \in [0, 1)$, there exists an $\alpha \geq 0$ such that the continuum discriminant direction w_γ is given by the ridge estimator w_α^R. This parallels the observation made by Sundberg (1993) in regression contexts. We allow negative α, so that the relation to ridge estimators is extended for $\gamma > 1$.

Theorem 1. *If d is not orthogonal to all eigenvectors corresponding to the largest eigenvalue λ_1 of S_T, then there exists a number $\alpha \in (-\infty, -\lambda_1) \cup [0, \infty)$ such that $w_\gamma \propto (S_T + \alpha I)^- d$, including the limiting cases $\alpha \to 0, \alpha \to \pm\infty$ and $\alpha \to -\lambda_1$.*

The above theorem can be shown by an application of Proposition 2.1 of Bjorkstrom and Sundberg (1999) who showed that, in our notation, the solution of $\max_w T_\gamma(w)$ is of the ridge form. See Appendix for detailed discussion on the proof of the theorem.

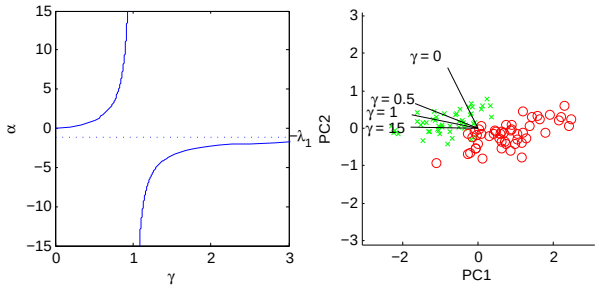

FIG 2. *Relation between γ and α, illustrated for the iris dataset, and w_γ for different γ values overlaid on the scatter plot of the first two principal components of the iris data. Different symbols represent different groups.*

The relation between α and γ is nonlinear and depends on S_T and d. A typical form of relation is plotted in Fig. 2, and is explained in the following example.

Example 2. From Fisher's iris data (Fisher, 1936), we chose 'versicolor' and 'virginica' as two groups each with 50 samples. For presentational purposes, we take the first two principal components of the dataset. See Fig. 2 for the scatterplot of the first two principal components. For a dense set of $\gamma \in [0, \infty)$, the corresponding α is plotted (in the left panel of Fig. 2), which exhibits the typical relationship between γ and α. The MDP at $\gamma = 0$ corresponds to the ridge solution with $\alpha = 0$. As γ approaches 1, the corresponding ridge solution is obtained with $\alpha \to \pm\infty$. For $\gamma > 1$, α is negative and approaches $-\lambda_1$ as $\gamma \to \infty$. The continuum directions $\{w_\gamma : \gamma \in [0, \infty)\}$ range from w_{LDA} (which is the same as w_{MDP} since $n > p$) to w_{PCA} (See the right panel of Fig. 2).

The ridge solution may not give a global maximum of T_γ when the assumption in Theorem 1 does not hold. An analytic solution for such a case is also provided in Proposition 6 in the Appendix.

2.4. Continuum directions in high dimensions

In high dimensional situations where the dimension p of the primary data is much higher than the sample size n, the continuum directions are still well-defined. We return to discuss that MDP has more preferable properties than LDA in the high dimension low sample size situation. The ridge solution plays an important role in the following discussion.

In the conventional case where $p \leq n - 2$, It is easy to see that the ridge criterion (3) and its solution w_α^R bridge LDA and MD. However in high dimension low sample size situations (that is, S_W is rank deficient), one extreme of the ridge criterion is connected to MDP but not to LDA. The following proposition

shows that the parameter α in (3) bridges MDP and MD, giving a reason to favor MDP over LDA.

Proposition 2. *For $\alpha > 0$, $w_\alpha^R \propto (S_T + \alpha I)^{-1} d$. Moreover w_α^R is continuous with respect to $\alpha \in (0, \infty)$. The boundaries meet MDP and MD directions, that is, $\lim_{\alpha \to 0} w_\alpha^R = w_{MDP}$ and $\lim_{\alpha \to \infty} w_\alpha^R = w_{MD}$.*

Proof of Proposition 2. We first show that $(S_T + \alpha I)^{-1} d \propto (S_W + \alpha I)^{-1} d$. Let $\Omega = S_W + \alpha I$, whose inverse exists for $\alpha > 0$. Then $S_T + \alpha I = \Omega + c_0 dd^T$ for $c_0 = \frac{n_1 n_2}{n^2}$. By Woodbury's formula, $(S_T + \alpha I)^{-1} = \Omega^{-1} - c_1 \Omega^{-1} dd^T \Omega^{-1}$ for some constant c_1. Therefore, $(S_T + \alpha I)^{-1} d = \Omega^{-1} d - c_1 \Omega^{-1} dd^T \Omega^{-1} d = c_2 \Omega^{-1} d \propto (S_W + \alpha I)^{-1} d$.

The ridge solution w_α^R lies in the range of S_T, as shown in Lemma 4 in the Appendix. Writing w_α^R in the eigen-coordinates of S_T makes the proof simple. Let $S_T = U \Lambda U^T$ be the eigen-decomposition of S_T with $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_m)$, $U = [u_1, \ldots, u_m]$ for $m = \min(n-1, p)$. Then for $z_\alpha^R = U^T w_\alpha^R$ and $\delta = (\delta_1, \ldots, \delta_m)^T = U^T d$, we have $z_\alpha^R \propto (\frac{\delta_1}{\lambda_1 + \alpha}, \ldots, \frac{\delta_m}{\lambda_m + \alpha})^T$, which leads to the continuity of $w_\alpha^R = U z_\alpha^R$ with respect to $\alpha \in [0, \infty)$. It is now easy to see that $w_\alpha^R \to w_{MDP} \propto U \Lambda^{-1} \delta$ as $\alpha \to 0$. For the last argument, $z_\alpha^R \propto \alpha (\frac{\delta_1}{\lambda_1 + \alpha}, \ldots, \frac{\delta_m}{\lambda_m + \alpha})^T \to \delta$ as $\alpha \to \infty$. □

While w_{MDP} is a limit of ridge solutions, w_{LDA} may not meet with w_α^R as α decreases to 0. When $p \gg n$, w_{MDP} is orthogonal to w_{LDA} if the mean difference d is not in the range of S_W, i.e., $\mathrm{rank}(S_W) < \mathrm{rank}(S_T)$ (Ahn and Marron, 2010),. This fact and Proposition 2 give $\lim_{\alpha \to 0} \mathrm{angle}(w_{LDA}, w_\alpha^R) = 90°$.

Algebraically, the discontinuity of the ridge direction to w_{LDA} comes from the discontinuity of the pseudoinverse. Heuristically, the discontinuity comes from the fact that d does not completely lie in the range of S_W. In such a case, there is a direction vector w_0 orthogonal to the range of S_W containing information about d (i.e., $d^T w_0 \neq 0$). Using S_W^- in LDA ignores such information. On the other hand, MDP uses S_T^-, which preserves all information contained in w_0.

The values of Fisher's criterion for various w's in Fig. 3 exemplifies that w_{MDP} should be used as Fisher discriminant direction rather than w_{LDA} in high dimensions. In our application of the continuum directions to binary discriminant analysis in Section 3, we check that the empirical performance of LDA is among the worst.

2.5. Numerical algorithms for general case

In general cases where the rank of the between-variance matrix S_B is greater than 1, the connection to generalized ridge solutions in Theorem 1 does not hold. Even with binary supervision, when a sequence of continuum directions $\{w_{(1)}, \ldots, w_{(\kappa)}\}$ is desirable, the ridge parameter $\alpha(\gamma)$ is different for different k in $w_{(k)}$, even when γ is held fixed. Here, we propose a gradient descent algorithm to sequentially solve (7) for a given γ.

We first discuss a gradient descent algorithm for $w_{(1)}$. Since the only constraint is that the vector w is of unit size, the unit sphere $S^{p-1} = \{w \in \mathbb{R}^p :$

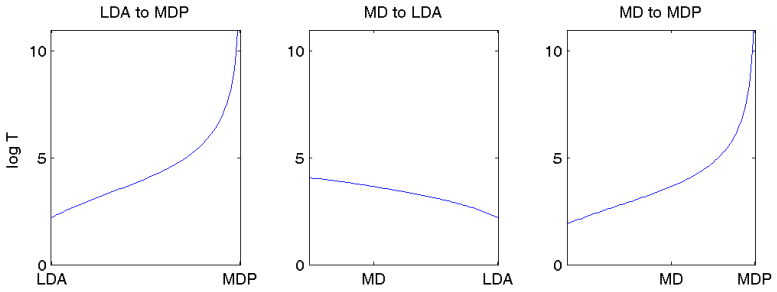

FIG 3. *Fisher's $T(w)$ for directions discriminating two groups ($n_1 = 20, n_2 = 17$) in a microarray dataset with $p = 2530$ (Bhattacharjee et al., 2001). The three horizontal axes represent discriminant direction w along the edges of the triangle formed by w_{LDA}, w_{MD}, and w_{MDP}. LDA is not maximizing Fisher's criterion and is inferior to the mean difference, while $T(w_{MDP}) = \infty$.*

$w^T w = 1\}$ is the feasible space. To make the iterate confined in the feasible space we update a candidate w_0 with $w_1 = (w_0 + c\nabla_{w_0})/\|w_0 + c\nabla_{w_0}\|$, for a step size $c > 0$, where the gradient vector is $\nabla_w = \frac{S_B w}{w^T S_B w} + (\gamma - 1)\frac{S_T w}{w^T S_T w}$. To expedite convergence, c is initially chosen to be large so that $w_1 \approx \nabla_{w_0}/\|\nabla_{w_0}\|$. If this choice of c overshoots, i.e., with this update $T_\gamma(w_1) < T_\gamma(w_0)$ then we immediately reduce c to unity, so that the convergence to maximum is guaranteed, sacrificing fast rate of convergence. The iteration is stopped if $1 - |w_1^T w_0| < \varepsilon$ or $|T_\gamma(w_1) - T_\gamma(w_0)| < \varepsilon$ for a needed precision $\varepsilon > 0$. The step size c can be reduced if needed, setting $c = 1$ has ensured convergence with a precision level $\varepsilon = 10^{-10}$ in our experience.

For the second and subsequent directions, suppose we have $w_{(1)}, \ldots, w_{(k)}$ and are in search for the $w_{(k+1)}$. The S_T-orthogonality and unit size constraints lead to the feasible space $\mathcal{S} = \{w \in S^p : w^T S_T w_{(\ell)} = 0, w^T w = 1, \ell = 1, \ldots, k\}$. Since any $w \in \mathcal{S}$ is orthogonal to $z_{(\ell)} := S^T w_{(\ell)}$, $\ell = 1, \ldots, k$, the solution lies in the nullspace of $Z_k = [z_{(1)}, \ldots, z_{(k)}]$. We use orthogonal projection matrix $P_k = I - Z_k(Z_k^T Z_k)^{-1} Z_k$ to project the variance-covariance matrices S_T and S_B onto the nullspace of Z_k, and obtain $S_T^{(k)} = P_k S_T P_k$ and $S_B^{(k)} = P_k S_B P_k$. The gradient descent algorithm discussed above for $w_{(1)}$ is now applied with $S_B^{(k)}$ and $S_T^{(k)}$ to update candidates of $w_{(k+1)}$, without the S_T-orthogonality constraint.

The following lemma justifies this iterative algorithm converges to the solution $w_{(k+1)}$.

Lemma 3. *(i) Let $x_i^* = P_k x_i$ be the projection of x_i onto the nullspace of Z_k. Write $X^* = [x_1^*, \ldots, x_n^*]$. Then $S_T^{(k)} = n^{-1} X^* (X^*)^T$ and $S_B^{(k)} = n^{-1}(X^* Y^T)(X^* Y^T)^T$.*

(ii) For $w \in \mathcal{S}$, $T_\gamma(w) = (w^T S_B^{(k)} w)(w^T S_T^{(k)} w)^{\gamma - 1} := T_\gamma^{(k)}(w)$.

(iii) The solution $w_{(k+1)}$ of the unconstrained optimization problem $max_w T_\gamma^{(k)}$

satisfies $w_{(k+1)}^T S_T w_\ell = 0$ *for* $\ell = 1, \ldots, k$.

Proof of Lemma 3. Part (i) is trivial. For part (ii), note that for all $w \in \mathcal{S}$, $P_k w = w$. Replacing w by $P_k w$ in $T_\gamma(w)$ gives the result. For part (iii), we use Lemma 4 in the Appendix which shows that the solution w of maximizing $T_\gamma^{(k)}$ lies in the column space of $P_k S_T P_k$. Thus, the solution $w_{(k+1)}$ satisfies the constraint $w_{(k+1)}^T S_T w_\ell = 0$ for $\ell = 1, \ldots, k$. □

It can be seen from Lemma 3 that the optimization is performed with the part of data that is S_T-orthogonal to Z_k. While making the optimization simpler, we do not lose generality because the original criterion T_γ has the same value as $T_\gamma^{(k)}$ for candidate w in the feasible region (Lemma 3(ii)). This with the last result (iii) shows that our optimization procedure leads to (at least) local maximum in the feasible region.

3. Application to binary classification

In this section, we focus on an application of the continuum directions to linear binary classification. In particular, we return to the simple setting presented in Section 2.1, and compute w_γ as in (6), which is normal to a separating hyperplain. Suppose that w_γ points towards the first group, i.e., $w_\gamma^T d > 0$. A new observation x is classified to the first group if $(x - \hat{\mu})^T w_\gamma > -\log(n_1/n_2)$. Since the direction w_γ is directly used in classification, we call it a continuum discriminant direction and the classifier will be called continuum discriminant analysis (CDA), in this section.

3.1. Computation and cross-validation for CDA

A discrete sequence of the first continuum directions is obtained by the relationship between w_γ and ridge solutions; See Section 2.3. We utilize the fact that there is a corresponding γ for each ridge parameter $\alpha \in (-\infty, -\lambda_1] \cup [0, \infty)$. Let $M > 0$ be a maximum value for evaluating α. In our experience it is sufficient to choose $M = 10\lambda_1$, ten times larger than the first eigenvalue of S_T. Define $\alpha_{(k)} = \frac{k}{K} M$ and $\alpha^{(k)} = -(1+\epsilon)\lambda_1 - \frac{K-k}{K} M$ for $k = 0, \ldots, K$ for some number K. The small number $\epsilon > 0$ keeps the matrix $S_T + \alpha^{(k)} I_p$ invertible and was chosen to 0.01 for numerical stability. For each $\alpha = \alpha_{(k)}$ or $\alpha^{(k)}$, we get $w_{\gamma(\alpha)} = (S_T + \alpha I_p)^{-1} d$, where

$$\gamma(\alpha) = \frac{\alpha}{w_{\gamma(\alpha)}^T S_T w_{\gamma(\alpha)} + \alpha}.$$

The sequence $\{w_{\gamma(\alpha)} : \alpha = \alpha_{(k)}, \alpha^{(k)}, k = 0, \ldots, K\}$ is augmented by the two extremes $w_{MD} (\propto d)$ and w_{PC1}.

If d is orthogonal to all eigenvectors corresponding to λ_1, then γ does not tend to infinity even though α has reached $-\lambda_1$. In such a case, the remaining sequence of directions is analytically computed using Proposition 6.

In the high dimension low sample size situation, where $p \gg n$, a computational shortcut to avoid the inversion of the $p \times p$ matrix $S_T + \alpha I_p$ can be used. Recall the eigen-decomposition of $S_T = U\Lambda U^T$ where $U = [u_1, \ldots, u_m]$ spans the range (column space) of S_T for $m = \min(n-1, p)$. Then for $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_m)$, $\delta = U^T d$, the continuum directions are obtained via only involving the inversion of $m \times m$ matrices:

$$(S_T + \alpha I_p)^- d = U(\Lambda + \alpha I_p)^{-1}\delta.$$

For the two-group classification problem at hand, we use the following cross validation scheme to choose the value of γ (or equivalently, α). Randomly divide the numbers $\{1, \ldots, n_1, n_1+1, \ldots, n\}$ into $2 \le R \le n$ non-empty subgroups $H_1, H_2, \ldots, H_R$. Then the samples are divided into $\mathbb{X}_r = \{x_i, y_j : i \in H_r, j+n_1 \in H_r\}$, for $r = 1, \ldots, R$. Let $S_T^{(r)}$, $\hat{\mu}^{(r)}$ and $d^{(r)}$ be defined as in Section 2.1 based on $\{x_1, \ldots, x_{n_1}, y_1, \ldots, y_{n_2}\} \backslash \mathbb{X}_r$. For any choice of γ, $w_\gamma^{(r)}$ is computed based on $S_T^{(r)}$ and $d^{(r)}$ as explained earlier in this section. For $j \in H_r$, let $M_j^{(r)} = 0$ if $j \le n_1$ (x_j is the group 1) and $(x_j - \hat{\mu}^{(r)})^T w_\gamma^{(r)} > -\log(n_1/n_2)$, or if $j > n_1$ and $(y_{j-n_1} - \hat{\mu}^{(r)})^T w_\gamma^{(r)} < -\log(n_1/n_2)$. Then

$$CV(\gamma) = \sum_{r=1}^{R} \sum_{j \in H_r} M_j^{(r)}$$

is the total number of misclassified cases among the validation sets for the given γ. The final choice of γ is $\hat{\gamma} = \arg\min_\gamma CV(\gamma)$. If the minimum is attained at several γ's, the smallest of these γ's is selected.

The computational procedure is further demonstrated in Section 3.3.

3.2. A simulation study

A simulation is performed to empirically reveal the underlying model under which the continuum directions are useful and the numerical performance of CDA, the classifier based on the first continuum direction, is good. We compare with linear discriminant analysis (LDA), independence rule (IR) by Bickel and Levina (2004), the features annealed independence rule (FAIR) by Fan and Fan (2008), and the distance weighted discrimination (DWD) (Marron, Todd and Ahn, 2007; Qiao et al., 2009).

The setup in the simulation study is as follows. We assume two groups with mean $\mu_1 = 0$ and $\mu_2 = c_0 1_s$ for some constant c_0 and $1_s = (1, \ldots, 1, 0, \ldots, 0)^T$ where the number of 1's is $s = 10$. The common covariance matrix is $\Sigma_p = (1-\rho)I_p + \rho 1_p 1_p^T$ for $\rho \in \{0, 0.1, 0.25\}$. The compound symmetry model allows examination from independent to highly correlated settings. The scalar $c_0 = 3(1_s^T \Sigma_p^{-1} 1_s)^{-1/2}$ varies for different (p, ρ) to keep the Mahalanobis distance between μ_1 and μ_2 equal to 3.

Training and testing data of size $n_1 = n_2 = 50$ are generated from normal distribution of dimension $p = 200, 400$ and 800. The parameter γ of CDA is chosen

	p	CDA	LDA	IR	FAIR	DWD
$\rho = 0$	200	14.32 (3.45)	29.59 (5.31)	13.66 (3.18)	8.90 (3.10)	13.88 (3.33)
	400	19.70 (4.07)	34.76 (5.33)	19.28 (4.10)	9.02 (3.23)	19.28 (4.10)
	800	24.90 (4.78)	39.80 (4.97)	24.50 (4.79)	9.80 (3.46)	24.14 (4.36)
$\rho = 0.1$	200	11.27 (3.56)	20.37 (4.65)	49.97 (5.25)	48.25 (7.09)	36.99 (7.07)
	400	9.87 (3.30)	26.97 (6.04)	50.02 (5.17)	49.39 (5.09)	45.11 (5.11)
	800	12.94 (3.79)	36.24 (5.72)	50.50 (4.13)	50.32 (4.19)	48.65 (4.69)
$\rho = 0.25$	200	5.90(2.72)	13.38 (4.16)	49.34 (5.14)	48.98 (5.30)	42.37 (5.58)
	400	3.88 (2.15)	19.93 (4.61)	50.90 (5.00)	50.55 (5.20)	47.61 (5.20)
	800	5.67 (2.62)	31.14 (5.32)	49.19 (4.69)	49.05 (4.71)	48.03 (4.74)

TABLE 1

Compound Symmetry model with high dimension, low sample size data: Mean misclassification error (in percent) with standard deviation in parentheses.

by the ten-fold cross validation as discussed in the previous section. The number of features for FAIR is also chosen by the ten-fold cross validation. The average classification errors and the standard deviations based on 100 replications are provided in Table 1.

Our results show that the CDA performs much better than other methods considered here when the correlation between variables is strong ($\rho = 0.1, 0.25$). In the independent setting ($\rho = 0$), the performance of CDA is comparable to IR and DWD. FAIR is significantly better than CDA when $\rho = 0$, because the crucial assumption of FAIR that the non-zero coordinates of $\mu_1 - \mu_0$ are sparse is also satisfied. The IR and FAIR, based on the assumption of independent variables, perform well when $\rho = 0$, but severely suffer from larger correlations ($\rho = 0.1, 0.25$), in which cases classification rates are close to 50% which is about the same as random guess. DWD also suffers from the highly correlated structure.

Another observation is that the performance of LDA is better for larger ρs. A possible explanation is that the underlying distribution $N(\mu_i, \Sigma)$ becomes nearly degenerate for larger ρ. The true covariance matrix has a very large first eigenvalue $\lambda_1 = p\rho + (1 - \rho)$ compared to the rest of eigenvalues $\lambda_j = 1 - \rho$, $2 \leq j \leq p$. Both the LDA and CDA benefit from extensively incorporating the covariance structure, in spite of the poor estimation of Σ_p when $p \gg n$.

Poor performance of FAIR for the strongly correlated case is also reported in Fan, Feng and Tong (2011), where they proposed the regularized optimal affine discriminant (ROAD), which is computed by a coordinate descent algorithm. Due to the heavy computational cost, we excluded the ROAD as well as the sparse LDA by Wu et al. (2009) and the linear programming discriminant rule (LPD) by Cai and Liu (2011). These methods aim to select few features as well as to classify, based on assumptions of sparse signals. CDA does not require such assumptions.

3.3. Real Data Examples

In this section, we provide two real data examples whose auxiliary information is binary.

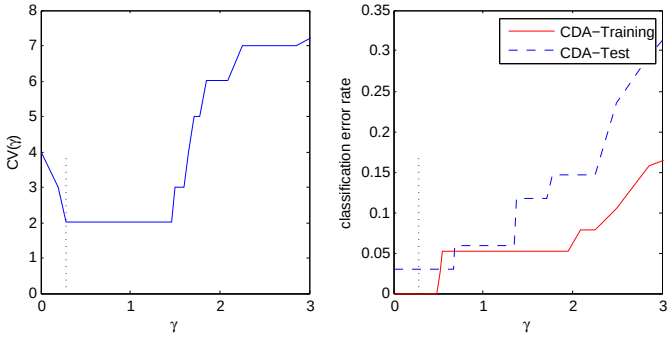

FIG 4. *Left: Cross validatory errors for* $\gamma \in [0, 3]$ *evaluated for Leukemia data. The* $\hat{\gamma} = 0.279$ *(located at the vertical dotted line) is the smallest* γ *that minimizes* $CV(\gamma)$*. Right: Classification error rates of training and testing set for different* γ*s.*

	CDA	LDA	IR	DWD	SVM	FAIR	LPD
Training error	0/38	1/38	1/38	0/38	0/38	1/38	0/38
Testing error	1/34	6/34	6/34	2/34	5/34	1/34	1/34

TABLE 2
Classification error of Leukemia data

3.3.1. Leukemia data

We first use the well-known data set of Golub et al. (1999), which consists of expression levels of the 7129 genes from 72 acute leukemia patients. Similar to the data preparation of Cai and Liu (2011), the dataset was preprocessed as follows. We filtered 140 genes with extreme variances, i.e., either larger than 10^7 or smaller than 10^3. Then genes with the 3000 largest absolute t-statistics were chosen. The dataset included 38 training cases (27 AMLs and 11 ALLs) and 34 testing cases (20 AMLs and 14 ALLs) and is available at `http://www.broadinstitute.org`.

With the two-group classification in mind, we obtain w_γ for a discrete set of $0 \leq \gamma < \infty$, as in Section 3.1. The cross validation with $R = 10$ leads to $\hat{\gamma} = 0.279$. Figure 4 shows the cross validation index $CV(\gamma)$ over $\gamma \in [0, 4]$. The smallest cross validatory misclassification rate is 2/38. Figure 4 also shows the classification errors of training and testing data for different γs. For smaller γ values, including $\gamma = 0$ (corresponding to MDP) and the chosen $\hat{\gamma}$, the classification errors are 1 out 34 for the test set, and 0 out of 38 for the training set. In comparison, LDA, the independence rule, DWD and linear SVM (Vapnik, 1995) result in 2–6 testing errors. From the work of Fan and Fan (2008) and Cai and Liu (2011), FAIR and LPD makes only 1/34 testing error. The classification results of CDA with the chosen $\hat{\gamma}$ as well as those of other methods are summarized in Table 2.

	CDA	LDA	DWD	FAIR	QDA
Test	33.7 (6.38)	37.4 (6.48)	33.6 (6.33)	33.3 (6.17)	34.4 (6.85)
Train	33.3 (0.79)	13.9 (1.03)	30.7 (0.78)	32.6 (0.84)	6.7 (2.85)

TABLE 3

Misclassification rate (in percent) of liver nuclei outlines data. Mean and standard deviation of ten repetitions are reported.

3.3.2. Liver cell nuclei shapes

In a biomedical study, it is of interest to quantify the difference between normal and cancerous cell nuclei, based on the shape of cells. We analyze a dataset in Wang et al. (2011a), which consists of discretized cell outlines, Procrustes-fitted to each other to extract shape information. See Rohde et al. (2008); Wang, Ozolek and Rohde (2010); Wang et al. (2011a,b) for detailed pre-processing procedures. In the dataset, there are $n_1 = 250$ normal liver tissue outlines and $n_2 = 250$ liver hepatoblastoma outlines. Each outline is represented by 90 planar landmarks, leading to $p = 180$.

In the context of classification of the above dataset, CDA with cross-validated γ is compared with LDA, DWD, FAIR, and a quadratic discriminant analysis (QDA). As explained in Section 3.2, the threshold value of FAIR is chosen using cross validation. The QDA is modified to have smaller variability by using a ridge-type covariance estimator. Explicitly, let S_1, S_2 be the sample covariance matrices for each group. Then the modified QDA classifies a new observation x to group 1 if $\delta_1(x) > \delta_2(x)$, where $\delta_k(x) = -1/2 \log[\det(S_k + \alpha I)] - 1/2(x - \hat{\mu}_k)^T (S_k + \alpha I)^{-1}(x - \hat{\mu}_k) + \log(n_k/n)$, for $k = 1, 2$. The ridge parameter α is chosen by a cross-validation similar to the procedure described in Section 3.1.

For the comparison, we randomly assign 50 cases as a testing data set, and each classifier is calculated with the remaining 450 cases. The empirical misclassification rates of classifiers are computed based on the training dataset and on the testing dataset. This is repeated for 100 times to observe the variation of the misclassification rates. For the continuum directions with varying γ, we observe that the classification error rates become stable as γ increases. See Fig. 5. Both the training and testing error rates become close to 33.5% as the directions tend to MD and to PCA. This is because, for this dataset, w_{MD} and w_{PC1} are close to each other with angle$(w_{MD}, w_{PC1}) = 6.67°$, and both exhibit good classification performances, with error rates close to 33.5%. For each training dataset, $\hat{\gamma}$ is chosen by the cross validation. Many $\hat{\gamma}$s have values between $(0, 0.5)$, but a few of those are as large as $\gamma = 3$, as shown in Fig. 5. The performance of CDA with cross-validated γ is compared with other methods in Table 3. Based on the testing error rate, CDA performs comparable to more sophisticated methods such as FAIR and DWD while LDA and QDA tend to overfit and result in worse classification rates than the other methods examined.

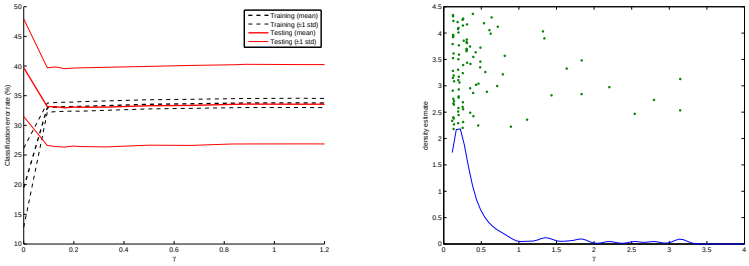

FIG 5. *Left: Classification error rates of training and testing set for different γs. Right: A jitter plot with a density estimate for values of $\hat{\gamma}$ chosen by the cross validation.*

4. Discussion

We proposed a criterion evaluating continuum directions, forming a basis for dimension reduction of primary data set, while the degrees of supervision from a auxiliary data set are controlled by a meta-parameter. An application of the proposed directions to binary classification was discussed in more detail. Numerical properties of the proposed binary classifier have demonstrated good performance when $p \gg n$. In particular, our method outperforms several other methods when the variance of the first principal component is much larger than the rest. We also check that in real data sets, our method performs similar to much complicated methods.

The proposed method bridge several well-known approaches, including LDA, MDP, MD, ridge estimators and PCA. The method provides a simple but unified framework in understanding the aforementioned methods, as continuum regression does in the regression setting. There are several other criteria that also give a transition between LDA (or MDP) and PCA. A slightly modified criterion from (6), $F_\alpha(w) = (w^T S_B w)^2/|w^T(S_T + \alpha I_p)w|$ with the constraint $w^T w = 1$, gives the ridge solution $\tilde{w}_\alpha = (S_T + \alpha I)^- d$ with the same $\alpha \in (-\infty, \lambda_1) \cup [0, \infty)$. This criterion is first introduced in a regression problem (Frank and Friedman, 1993; Bjorkstrom and Sundberg, 1999), but has not been adopted into classification framework. Wang et al. (2011b) proposed a modified Fisher's criterion

$$\tau_\delta(w) = \frac{w^T S_T w}{w^T(S_W + \delta I)w}, \tag{9}$$

that bridges between LDA and PCA. For $\delta = 0$, the criterion (9) becomes identical to equation (1) up to the constant 1, thus equivalent to LDA. In the limit of $\delta \to \infty$, $\delta\tau_\delta(w)$ converges to the criterion for w_{PC1}. The maximizer of τ_δ is a solution of a generalized eigenvalue problem. We leave further investigation of these criteria as future research directions.

Lee, Ahn and Jeon (2013) also discussed discrimination methods that bridge MDP and MD, in high dimension low sample size context. The method of Lee,

Ahn and Jeon (2013) is in fact equivalent to a part of continuum directions, restricted for $\gamma \in [0, 1]$. In this paper, the continuum between MDP to PCA is completed by also considering $\gamma > 1$, the method is extended for supervised dimension reduction, and a connection to continuum regression is made clear.

Acknowledgments

The author is grateful to Gustavo Rohde for sharing the nuclei cell outlines data, and to Jeongyoun Ahn and Myung Hee Lee for their helpful suggestions and discussions.

Appendix

Continuum directions lie in the range of S_T.

Lemma 4. *Any maximizer w of (7) lies in the range (column space) of S_T.*

Proof of Lemma 4. Denote the range of S_T by $\mathcal{R}_T$. Let $\text{rank}(S_T) = m \leq \min(n-1, p)$. Then for any $w \in \mathbb{R}^p$ with $\|w\| = 1$, let w_P be the orthogonal projection of w onto $\mathcal{R}_T$. Then $\|w_P\| \leq 1$ where the equality holds if and only if $w \in \mathcal{R}_T$. Let $\tilde{w} = w_P/\|w_P\|$. Then since $w^T S_T w = w_P^T S_T w_P$ and $w^T S_B w = w_P^T S_B w_P$, we have for $\gamma \geq 0$,

$$T_\gamma(w) = (\tilde{w}^T S_B \tilde{w})(\tilde{w}^T S_B \tilde{w})^{\gamma-1}\|w_P\|^{2\gamma} \leq T_\gamma(\tilde{w}).$$

Thus the maximizer of $T_\gamma(w)$ always lies in $\mathcal{R}_T$. □

Proof of Theorem 1.

In a multivariate linear regression problem, with the $n \times p$ design matrix X and the n vector y of responses, denote a regressor by $w^T x$ a linear combination of p variables. Both X and y are assumed centered. Let $V(w) = w' X^T X w$ be the sample variance of the regressor. Let $K(w) = y^T X w$ be the sample covariance between the regressor and y and $R(w)$ be the sample correlation, which is proportional to $K/\sqrt{V}$. The following theorem is from Bjorkstrom and Sundberg (1999).

Theorem 5 (Proposition 2.1 of Bjorkstrom and Sundberg (1999))**.** *If a regressor w_f is defined according to the rule*

$$wf = arg \max_{\|w\|=1} f(K^2(w), V(w)),$$

where $f(K^2, V)$ is increasing in K^2 (or R^2) for constant V, and increasing in V for constant R^2, and if $X^T y$ is not orthogonal to all eigenvectors corresponding to the largest eigenvalue λ_1 of $X^T X$, then there exists a number α such that $w_f \propto (X^T X + \alpha I)^{-1}$, including the limiting cases $\alpha \downarrow 0, \alpha \uparrow \infty$ and $\delta \uparrow -\lambda_1$.

A two-group classification problem is understood as a special case of regression. In particular, let y be $+1$ if the ith observation is in the first group or -1 if it is in the second group. Then the total variance matrix $S_T \propto X^T X$ and the mean difference $d = X^T y$. The criterion (6) is $K^2(w)V^{\gamma-1}(w)$, which satisfies the assumptions of Theorem 5. Theorem 1 is thus a special case of Theorem 5.

Analytic solution for the rare case

The ridge solution may not give a global maximum of T_γ when the assumption in Theorem 1 does not hold. We give an analytic solution for such a case. It is convenient to write w in the canonical coordinates of S_T. Let $S_T = U\Lambda U^T$ be the eigen-decomposition of S_T with $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_m)$, $U = [u_1, \ldots, u_m]$ for $m = \min(n-1, p)$, with convention $\lambda_i \geq \lambda_{i+1}$. To incorporate any duplicity of the first eigenvalue let ι represent the number of eigenvalues having the same value as λ_1, that is, $\lambda_1 = \ldots = \lambda_\iota$. Denote $\Lambda_1 = \mathrm{diag}(\lambda_1, \ldots, \lambda_\iota) = \lambda_1 I_\iota$ and $\Lambda_2 = \mathrm{diag}(\lambda_{\iota+1}, \ldots, \lambda_m)$. Let $z = U^T w$ and $\delta = (\delta_1, \ldots, \delta_m)^T = U^T d$. It is shown in the Appendix that the maximizer of (6) lies in the range of S_T.

Proposition 6. *Suppose d is orthogonal to all eigenvectors corresponding to λ_1 and is not orthogonal to all eigenvectors corresponding to $\lambda_{\iota+1}$. Let*

$$z_\alpha = \frac{(\Lambda_2 + \alpha I)^{-1}\delta_2}{\sqrt{\delta_2^T(\Lambda_2 + \alpha I)^{-2}\delta_2}} \quad \text{for} \quad \alpha \in (-\infty, -\lambda_1] \cup [0, \infty).$$

(i) If $z_{-\lambda_1}^T(\lambda_1 I - \Lambda_2)z_{-\lambda_1} \leq \lambda_1/\gamma$, then $w_\gamma = U\tilde{z}$, $\tilde{z}^T = [0_\iota^T, z_\alpha^T]$ for some $\alpha \in (-\infty, -\lambda_1] \cup [0, \infty)$.

(ii) If $z_{-\lambda_1}^T(\lambda_1 I - \Lambda_2)z_{-\lambda_1} > \lambda_1/\gamma$, then there exist multiple solutions $w_\gamma = U\hat{z}$, $\hat{z}^T = (\hat{z}_1^T, \hat{z}_2^T)$, of (6) satisfying

$$\hat{z}_1 \in \{z_1 \in \mathbb{R}^\iota : z_1^T z_1 = 1 - \lambda_1/(\gamma z_{-\lambda_1}^T(\lambda_1 I - \Lambda_2)z_{-\lambda_1})\}$$

and

$$\hat{z}_2 = \sqrt{\frac{\lambda_1}{\gamma}}\frac{(\Lambda_2 - \lambda_1 I)^{-1}\delta_2}{\sqrt{\delta_2^T(\lambda_1 I - \Lambda_2)^{-1}\delta_2}}.$$

Proof of Proposition 6. Recall $z = U^T w$ and $\delta = (\delta_1, \ldots, \delta_m)^T = U^T d$. Grouping z and δ into the first ι elements and the rest, write $z^T = (z_1^T, z_2^T)$, $\delta^T = (\delta_1^T, \delta_2^T)$. If d is orthogonal to all eigenvectors corresponding to λ_1, then $\delta_1 = 0$. Rewriting equation (8) in the eigen-coordinates gives two systems of equations

$$0 + (\gamma - 1)\frac{\Lambda_1 z_1}{z_1^T\Lambda_1 z_1 + z_2^T\Lambda_2 z_2} - \gamma z_1 = 0, \tag{10}$$

$$\frac{\delta_2}{z_2^T\delta_2} + (\gamma - 1)\frac{\Lambda_2 z_2}{z_1^T\Lambda_1 z_1 + z_2^T\Lambda_2 z_2} - \gamma z_2 = 0. \tag{11}$$

If $\|z_1\| > 0$, then we have from (10)

$$\lambda_1\|z_1\|^2 = z_1^T\Lambda_1 z_1 = \frac{\gamma - 1}{\gamma}\lambda_1 - z_2^T\Lambda_2 z_2. \tag{12}$$

Equations (11) and (12) lead to

$$z_2 = c_\gamma(\Lambda_2 - \lambda_1 I)^{-1}\delta_2,$$

where c_γ satisfies $c_\gamma^2 = -\lambda_1/(\gamma\delta_2^T(\Lambda_2 - \lambda_1 I)^{-1}\delta_2)$, which is obtained from the constraint $\|z_1\|^2 + \|z_2\|^2 = 1$. Finally, we check that such a solution exists if $z_2^T z_2 \leq 1$, that is,

$$\frac{\sum_{i=\iota+1}^m \lambda_1\delta_i^2/(\lambda_1 - \lambda_i)^2}{\sum_{i=\iota+1}^m \delta_i^2/(\lambda_1 - \lambda_i)} \leq \gamma. \tag{13}$$

The criterion T_γ in the canonical coordinate is proportional to

$$T_\gamma(z) = (z_2^T\delta_2)^2(\lambda_1 z_1^T z_1 + z_2^T\Lambda_2 z_2)^{\gamma-1}.$$

Thus T_γ is maximized by $\hat{z}^T = (\hat{z}_1^T, \hat{z}_1^T)$ for any $\hat{z}_2 = \pm z_2$ and any $\hat{z}_1$ that satisfies (12). This proves (ii).

If (13) does not hold, then by contradiction we have $\|z_1\| = 0$. Thus $\tilde{z}$ is of the form $(0_\iota, z_2)$ for z_2 satisfying (11). Since the first coordinate of δ_2 is nonzero, an application of Theorem 1 leads that there exists $\alpha \in (-\infty, -\lambda_{\iota+1}) \cup [0, \infty)$ such that $z_2 \propto (\Lambda_2 + \alpha I)^-\delta_2$.

To conclude (i), we need to rule out the possibility of α having values in $(-\lambda_1, -\lambda_{\iota+1})$. Let $M_k = M_k(a) = \delta_2^T(aI - \Lambda_2)^{-k}\delta_2$ for $k = 1, 2, \ldots$. The derivative of M_k with respect to a is $M_k' = -kM_{k+1}$. We have $M_k(a) > 0$ for $a \in (\lambda_{\iota+1}, \lambda_1]$. The assumption of (i) is written as $\gamma \leq \lambda_1 M_2(\lambda_1)/M_1(\lambda_1)$. It can be shown that $aM_2(a)/M_1(a)$ is a decreasing function of $a > \lambda_{\iota+1}$. This leads to

$$\gamma \leq aM_2/M_1, \quad \text{for any} \quad a \in (\lambda_{\iota+1}, \lambda_1]. \tag{14}$$

For $z_\alpha = (\Lambda_2+\alpha I)^-\delta_2/\|(\Lambda_2+\alpha I)^-\delta_2\|$, $T_\gamma((0_\iota, z_{-a})) = M_1^2/M_2(a-M_1/M_2)^{\gamma-1}$, and the derivative of $\log(T_\gamma)$

$$\frac{2(M_2^2 - M_1M_3)}{M_1M_2(M_2a - M_1)}(\gamma M_1 - M_2 a) \geq 0 \quad \text{for any} \quad a \in (\lambda_{\iota+1}, \lambda_1].$$

We have used (14) and the Cauchy-Schwartz inequality. Since T_γ is increasing in a, any z_α with $\alpha \in (-\lambda_1, -\lambda_{\iota+1})$ can not be a maximizer of T_γ for any γ, which completes the proof. □

Example 3. Let $S_T = \mathrm{diag}(4, 1)$ and $d = (0, 2)^T$. This is the case where d is orthogonal to the first eigenvector of S_T. The assumption of Theorem 1(ii) holds for $\gamma > 4/3$. Considering the sequence of w_γ, a ridge solution exists for each $\gamma < 4/3$. The corresponding α reaches to $-\lambda_1 = -4$ at $\gamma = 4/3$; For $\gamma > 4/3$, Theorem 1 gives the exact expression for w_γ. In the toy example, $w_{LDA} = w_{MDP} = w_{MD} = w_\gamma$ for $\gamma \in [0, 4/3]$. For $\gamma > 4/3$, w_γ approaches w_{PCA} as γ increases. See Fig. 6 for an illustration of the continuum discriminant directions for the toy data.

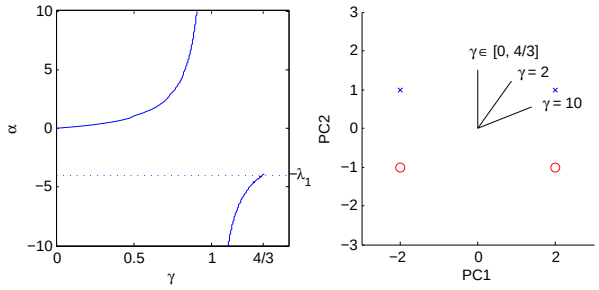

FIG 6. *Relation between γ and α, for the toy example. For $\gamma > 4/3$, w_γ does not correspond to any ridge solution.*

References

AHN, J. and MARRON, J. S. (2010). The maximal data piling direction for discrimination. *Biometrika* **97** 254–259.

BAIR, E., HASTIE, T., PAUL, D. and TIBSHIRANI, R. (2006). Prediction by Supervised Principal Components. *Journal of the American Statistical Association* **101** 119–137.

BARRETT, T., WILHITE, S. E., LEDOUX, P., EVANGELISTA, C., KIM, I. F., TOMASHEVSKY, M., MARSHALL, K. A., PHILLIPPY, K. H., SHERMAN, P. M., HOLKO, M. et al. (2013). NCBI GEO: archive for functional genomics data setsupdate. *Nucleic acids research* **41** D991–D995.

BHATTACHARJEE, A., RICHARDS, W. G., STAUNTON, J., LI, C., MONTI, S., VASA, P., LADD, C., BEHESHTI, J., BUENO, R., GILLETTE, M., LODA, M., WEBER, G., MARK, E. J., LANDER, E. S., WONG, W., JOHNSON, B. E., GOLUB, T. R., SUGARBAKER, D. J. and MEYERSON, M. (2001). Classification of human lung carcinomas by mRNA expression profiling reveals distinct adenocarcinoma subclasses. *Proc. Natl. Acad. Sci. USA* **98(24):137**.

BICKEL, P. J. and LEVINA, E. (2004). Some theory for Fisher's linear discriminant function , 'naive Bayes', and some alternatives when there are many more variables than observations. *Bernoulli* **10** 989–1010.

BJORKSTROM, A. and SUNDBERG, R. (1999). A Generalized View on Continuum Regression. *Scandinavian Journal of Statistics* **26** 17–30.

BOCK, C. (2012). Analysing and interpreting DNA methylation data. *Nature Reviews Genetics* **13** 705–719.

CAI, T. and LIU, W. (2011). A Direct Estimation Approach to Sparse Linear Discriminant Analysis. *Journal of the American Statistical Association* **106** 1566–1577.

CONNOR, G., HAGMANN, M. and LINTON, O. (2012). Efficient semiparametric estimation of the Fama–French model and extensions. *Econometrica* **80** 713–754.

Cook, R. D., Li, B. and Chiaromonte, F. (2010). Envelope models for parsimonious and efficient multivariate linear regression. *Statist. Sinica* **20** 927–1010.

Cook, R. D. and Ni, L. (2005). Sufficient dimension reduction via inverse regression. *Journal of the American Statistical Association* **100**.

Cook, R. D. and Su, Z. (2013). Scaled envelopes: scale-invariant and efficient estimation in multivariate linear regression. *Biometrika* **100** 939–954.

Cristianini, N. and Shawe-Taylor, J. (2000). *An Introduction to Support Vector Machines*. Cambridge University Press.

de Jong, S. and Farebrother, R. W. (1994). Extending the relationship between ridge regression and continuum regression. *Chemometrics and Intelligent Laboratory Systems* **25** 179–181.

Di, C.-Z., Crainiceanu, C. M., Caffo, B. S. and Punjabi, N. M. (2009). Multilevel functional principal component analysis. *The annals of applied statistics* **3** 458.

Fan, J. and Fan, Y. (2008). High Dimensional Classification Using Features Annealed Independence Rules. *Annals of statistics* **36** 2605–2637.

Fan, J., Feng, Y. and Tong, X. (2011). A ROAD to Classification in High Dimensional Space. *JRSSB*.

Fan, J., Liao, Y. and Wang, W. (2014). Projected Principal Component Analysis in Factor Models. *Available at SSRN 2450770*.

Fisher, R. A. (1936). The use of multiple measurements in taxonomic problems. *Annals of Eugenics* **7** 179–188.

Frank, I. E. and Friedman, J. H. (1993). A statistical view of some chemometrics regression tools. *Technometrics* **35** 109–135.

Gilad, Y., Rifkin, S. A. and Pritchard, J. K. (2008). Revealing the architecture of gene regulation: the promise of eQTL studies. *Trends in genetics* **24** 408–415.

Golub, T. R., Slonim, D. K., Tamayo, P., Huard, C., Gaasenbeek, M., Mesirov, J. P., Coller, H., Loh, M. L., Downing, J. R., Caligiuri, M. a., Bloomfield, C. D. and Lander, E. S. (1999). Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *Science (New York, N.Y.)* **286** 531–537.

Hastie, T., Tibshirani, R. and Friedman, J. (2009). *Elements of Statistical Learning*. Springer.

Hotelling, H. (1936). Relations between two sets of variates. *Biometrika* 321–377.

Izenman, A. J. (1975). Reduced-rank regression for the multivariate linear model. *Journal of multivariate analysis* **5** 248–264.

Jolliffee, I. T. (2002). *Principal Component Analysis*, 2nd ed. ed. *Springer Series in Statistics*. Springer-Verlag, New York.

Lee, M. H., Ahn, J. and Jeon, Y. (2013). HDLSS discrimination with adaptive data piling. *Journal of Computational and Graphical Statistics* **22** 433–451.

Lee, J., Dobbin, K. K. and Ahn, J. (2014). Covariance adjustment for batch effect in gene expression data. *Statistics in medicine* **33** 2681–2695.

Li, G., Yang, D., Nobel, A. B. and Shen, H. (2015a). Supervised singular value decomposition and its asymptotic properties. *To appear Journal of Multivariate Analysis*.

Li, G., Yang, D., Nobel, A. B. and Shen, H. (2015b). Supervised singular value decomposition and its asymptotic properties. *Journal of Multivariate Analysis*.

Liu, Y., Aryee, M. J., Padyukov, L., Fallin, M. D., Hesselberg, E., Runarsson, A., Reinius, L., Acevedo, N., Taub, M., Ronninger, M. et al. (2013). Epigenome-wide association data implicate DNA methylation as an intermediary of genetic risk in rheumatoid arthritis. *Nature biotechnology* **31** 142–147.

Lock, E. F., Hoadley, K. A., Marron, J. and Nobel, A. B. (2013). Joint and individual variation explained (JIVE) for integrated analysis of multiple data types. *The annals of applied statistics* **7** 523.

Marron, J. S., Todd, M. J. and Ahn, J. (2007). Distance Weighted Discrimination. *Journal of the American Statistical Association* **102** 1267–1271.

Mei, H., Hsieh, Y., Nardo, C., Xu, X., Wang, S., Ng, K. and Korfmacher, W. A. (2003). Investigation of matrix effects in bioanalytical high-performance liquid chromatography/tandem mass spectrometric assays: application to drug discovery. *Rapid Communications in Mass Spectrometry* **17** 97–103.

Qiao, X., Zhang, H. H., Todd, M. J. and Marron, J. S. (2009). Weighted Distance Weighted Discrimination and its Asymptotic Properties. *Operations Research* **105** 401–414.

Rohde, G. K., Ribeiro, A. J. S., Dahl, K. N. and Murphy, R. F. (2008). Deformation-based nuclear morphometry: capturing nuclear shape variation in HeLa cells. *Cytometry Part A the journal of the International Society for Analytical Cytology* **73** 341–50.

Shao, J., Wang, Y., Deng, X. and Wang, S. (2011). Sparse linear discriminant analysis by thresholding for high dimensional data. *The Annals of Statistics* **39** 1241–1265.

Shedden, K., Taylor, J. M., Enkemann, S. A., Tsao, M.-S., Yeatman, T. J., Gerald, W. L., Eschrich, S., Jurisica, I., Giordano, T. J., Misek, D. E. et al. (2008). Gene expression–based survival prediction in lung adenocarcinoma: a multi-site, blinded validation study. *Nature medicine* **14** 822–827.

Stone, M. and Brooks, R. J. (1990). Continuum Regression: Cross-Validated Sequentially Constructed Prediction Embracing Ordinary Least Squares, Partial Least Squares and Principal Components Regression. *Journal of the Royal Statistical Society. Series B (Methodological)* **52** 237–269.

Su, Z. and Cook, R. D. (2011). Partial envelopes for efficient estimation in multivariate linear regression. *Biometrika* asq063.

Su, Z. and Cook, R. D. (2012). Inner envelopes: efficient estimation in multivariate linear regression. *Biometrika* **99** 687–702.

Sundberg, R. (1993). Continuum Regression and Ridge Regression. *Journal of the Royal Statistical Society Series B Methodological* **55** 653–659.

Tso, M.-S. (1981). Reduced-rank regression and canonical analysis. *Journal of the Royal Statistical Society. Series B (Methodological)* 183–189.

Vapnik, V. N. (1995). *The Nature of Statistical Learning Theory. Statistics for Engineering and Information Science* **8**. Springer.

Wang, W., Ozolek, J. A. and Rohde, G. K. (2010). Detection and classification of thyroid follicular lesions based on nuclear structure from histopathology images. *Cytometry Part A the journal of the International Society for Analytical Cytology* **77** 485–494.

Wang, W., Ozolek, J. A., Slepčev, D., Lee, A. B., Chen, C. and Rohde, G. K. (2011a). An optimal transportation approach for nuclear structure-based pathology. *IEEE Transactions on Medical Imaging* **30** 621–631.

Wang, W., Mo, Y., Ozolek, J. A. and Rohde, G. K. (2011b). Penalized Fisher Discriminant Analysis and Its Application to Image-Based Morphometry. *Pattern Recognition Letters* **32** 2128–2135.

Witten, D. M. and Tibshirani, R. J. (2009). Extensions of sparse canonical correlation analysis with applications to genomic data. *Statistical applications in genetics and molecular biology* **8** 1–27.

Wold, S., Kettaneh, N. and Tjessem, K. (1996). Hierarchical multiblock PLS and PC models for easier model interpretation and as an alternative to variable selection. *Journal of chemometrics* **10** 463–482.

Wu, M. C., Zhang, L., Wang, Z., Christiani, D. C. and Lin, X. (2009). Sparse linear discriminant analysis for simultaneous testing for the significance of a gene set/pathway and gene selection. *Bioinformatics (Oxford, England)* **25** 1145–51.

Yang, E., Ravikumar, P., Allen, G. I., Baker, Y., Wan, Y.-W. and Liu, Z. (2014). A General Framework for Mixed Graphical Models. *arXiv preprint arXiv:1411.0288*.